arXiv:1510.06138v1 [stat.ML] 21 Oct 2015

Multiple co-clustering based on nonparametric mixture models with heterogeneous marginal distributions

Tomoki Tokuda *[1], Junichiro Yoshimoto[1], Yu Shimizu[1], Shigeru Toki[2], Go Okada[2], Masahiro Takamura[2], Tetsuya Yamamoto[3], Shinpei Yoshimura[4], Yasumasa Okamoto[2], Shigeto Yamawaki[2], and Kenji Doya[1]

[1]*Okinawa Institute of Science and Technology Graduate University, 1919-1 Tancha, Okinawa, 904-0495, JAPAN*

[2]*Department of Psychiatry and Neurosciences, Hiroshima University, Kasumi1-2-3, Minami-ku, Hiroshima, 734-8851, Hiroshima, JAPAN*

[3]*Japan Society for the promotion of science*

[4]*Otemon Gakuin University, Nishiai 2-1-15, Ibaraki city, Osaka, 567-8502, JAPAN*

Abstract

We propose a novel method for multiple clustering that assumes a co-clustering structure (partitions in both rows and columns of the data matrix) in each view. The new method is applicable to high-dimensional data. It is based on a nonparametric Bayesian approach in which the number of views and the number of feature-/subject clusters are inferred in a data-driven manner. We simultaneously model different distribution families, such as Gaussian, Poisson, and multinomial distributions in each cluster block. This makes our method applicable to datasets consisting of both numerical and categorical variables, which biomedical data typically do. Clustering solutions are based on variational inference with mean field approximation. We apply the proposed method to synthetic and real data, and show that our method outperforms other multiple clustering methods both in recovering true cluster structures and in computation time. Finally, we apply our method to a depression dataset with no true cluster structure available, from which useful inferences are drawn about possible clustering structures of the data.

Key-words: Clustering, multiple views, nonparametric mixture models, variational inference, high-dimensional data analysis

1 Introduction

We consider a clustering problem for a data matrix that consists of objects in rows and features (variables, or attributes) in columns. Clustering objects based on the data matrix is a basic data

*tomoki.tokuda@oist.jp

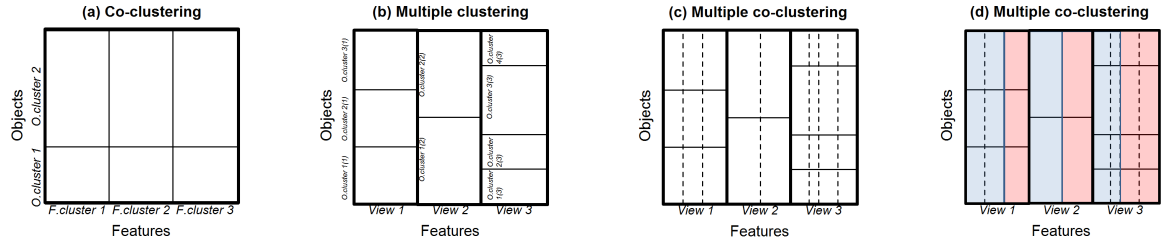

Figure 1: *Illustration of clustering structures: Panel (a) co-clustering; (b) multiple clustering (with full covariance of Gaussian); (c) multiple clustering with a specific structure of co-clustering; (d) extension of the model (c) where different distribution families are mixed (two distributions families in blue and red). Note that a rectangle surrounded by bold lines corresponds to a single co-clustering structure with a single object cluster solution. In these panels, features and objects are sorted in the order of view, feature and object cluster indices (hence, the order of objects differs among the co-clustering rectangles).*

mining approach, which groups objects with similar patterns of distribution. As an extension of conventional clustering, a co-clustering model has been proposed, which captures not only object cluster structure, but also feature cluster structure (features are grouped based on their distribution patterns, Figure 1a). This has the effect of reducing the number of parameters, which enables the model to fit high-dimensional data. Yet, the co-clustering method (as well as conventional clustering methods) often does not work well for high-dimensional data, because such data may have different 'views' that characterize multiple clustering solutions [1] [13, 16]. For instance, in DNA analysis, a set of genes (features) may be related to clustering of subjects for a specific genetic disorder, while another set may be related to clustering of subjects for a different disorder.

Recently, several heuristic methods have been proposed to detect multiple clustering solutions of objects [3, 1, 11, 4]. However, with these multiple clustering methods it is not straightforward to determine the number of views. A more promising approach based on nonparametric mixture models assumes multivariate Gaussian mixture models for each view (Figure 1b) [8]. In the latter approach, the full Gaussian model for covariance matrices is considered, and the numbers of views and of object clusters are inferred in a data-driven way via the Dirichlet process. Such a method is quite useful to discover possible multiple cluster solutions when these numbers are not known in advance, which is usually the case. However, its application is limited to low dimensional cases ($p < n$), because in high-dimensional cases, the number of objects to infer posterior distribution for the full covariance matrix of the Gaussian distribution may be insufficient, resulting in over-fitting. Moreover, this method also suffers the drawback that features need to belong to the same distribution family.

To address the aforementioned problems, we propose a novel multiple clustering method (referred to hereafter as the multiple co-clustering method). Our model is based on the following extension of the co-clustering model. First, we consider multiple views of co-clustering structure (Figure 1c), where a univariate distribution is fitted to each cluster block [17]. Second, for each cluster block, the proposed method simultaneously deals with an ensemble of several types of distribution families such as Gaussian, Poisson, and multinomial distribution (Figure 1d). Obviously,

[1] We also use a terminology of 'clustering', meaning the whole set of clusters in a view.

the first extension enables our model to fit high-dimensional data, while the second enables it to fit data that include different types of features (numerical and categorical).

One may consider a multiple clustering model by simply fitting a univariate distribution to each view (hereafter, we call it the 'restricted multiple clustering method'). However, such an approach has the drawback that it may replicate similar object cluster solutions for different views. For instance, features that discriminate among object clusters in the same manner would be allocated to different views, if these are negatively correlated or if they have different scales (hence, redundant views). As a consequence, it would not only complicate interpretation, but would also lose discriminative power relative to features. In the present paper, we retain this method for performance comparisons with our method.

2 Model

As in [8], our model is based on nonparametric mixture models using the Dirichlet process [6, 5]. However, unlike the conventional Dirichlet process, we employ a hierarchical structure, because in our model, the allocation of features is determined in two steps: the first allocation to a view, and the second to a feature cluster in that view. Moreover, we allow for mixing of several types of features, such as mixtures of Gaussian, Poisson, and categorical/multinomial distributions. Note that in this paper, we assume that types of features are pre-specified by the user, and do not draw inferences about them from data. In the following section, we formulate our method to capture these two aspects. To estimate model parameters, we rely on a variational Bayes EM (Expectation Maximization) algorithm, which provides (iterative) updating equations of relevant parameters. In general, determining whether these updating equations may be expressed in closed form is a subtle problem. However, this is the case in our model, which provides an efficient algorithm to estimate views and feature-/object cluster solutions.

2.1 Multiple clustering model

We assume that a data matrix $\boldsymbol{X}$ consists of M types of distribution families that are known in advance. We decompose $\boldsymbol{X} = \{\boldsymbol{X}^{(1)}, \ldots, \boldsymbol{X}^{(m)}, \ldots, \boldsymbol{X}^{(M)}\}$ with data size $n \times d^{(m)}$ for $\boldsymbol{X}^{(m)}$, where m is an indicator for a distribution family ($m = 1, \ldots, M$). Further, we denote the number of views as V (common to all distribution families), the number of feature clusters $G_v^{(m)}$ for view v and distribution family m, and the number of object clusters K_v for view v (common to all distribution families). Moreover, for simplicity of notation, we use $G^{(m)} = \max_v G_v^{(m)}$ and $K = \max_v K_v$ to denote the number of features and the number of clusters, allowing for empty clusters.

With this notation, for i.i.d. $d^{(m)}$-dimensional random vectors $\boldsymbol{X}_1^{(m)}, \ldots, \boldsymbol{X}_n^{(m)}$ for distribution family m, we consider a $d^{(m)} \times V \times G^{(m)}$ feature-partition tensor (3rd-order) $\boldsymbol{Y}^{(m)}$ in which $Y_{j,v,g}^{(m)} = 1$ if feature j of distribution family m belongs to feature cluster g in view v (0 otherwise). Combining this for different distribution families, we let $\boldsymbol{Y} = \{\boldsymbol{Y}^{(m)}\}_m$. Similarly, we consider a $n \times V \times K$ object-partition (3rd-order) tensor $\boldsymbol{Z}$ in which $Z_{i,v,k} = 1$ if object i belongs to object cluster k in view v. Note that feature j belongs to one of the views (i.e., $\sum_{v,g} Y_{j,v,g}^{(m)} = 1$)

while object i belongs to each view (i.e., $\sum_k Z^{(m)}_{i,v,k} = 1$). Further, $\boldsymbol{Z}$ is common to all distribution families, which implies that our model estimates subject cluster solutions using information on all distribution families.

For a prior generative model of $\boldsymbol{Y}$, we consider a hierarchical structure of views and feature clusters: views are first generated, followed by generation of feature clusters. Thus, features are partitioned in terms of pairs of view and feature cluster memberships, which implies that the allocation of feature is jointly determined by its view and feature cluster. On the other hand, objects are partitioned into object clusters in each view, hence, we consider just a single structure of object clusters for $\boldsymbol{Z}$. We assume that these generative models are all based on a stick-breaking process as follows.

Generative model for feature clusters $\boldsymbol{Y}$

We let $\boldsymbol{Y}^{(m)}_{j\cdot\cdot}$ denote a view/feature cluster membership vector for feature j of distribution family m, which is generated by a hierarchical stick-breaking process:

$$
\begin{aligned}
w_v &\sim \text{Beta}(\cdot|1,\alpha_1),\ v = 1,2,\ldots \\
\pi_v &= w_v \prod_{t=1}^{v-1}(1-w_t), \\
w'^{(m)}_{g,v} &\sim \text{Beta}(\cdot|1,\alpha_2),\ g = 1,2,\ldots,m = 1,\ldots,M \\
\pi'^{(m)}_{g,v} &= w'^{(m)}_{g,v} \prod_{t=1}^{g-1}(1-w'^{(m)}_{t,v}), \\
\tau^{(m)}_{g,v} &= \pi_v \pi'^{(m)}_{g,v} \\
\boldsymbol{Y}^{(m)}_{j\cdot\cdot} &\sim \text{Mul}(\cdot|\boldsymbol{\tau}^{(m)}),
\end{aligned}
$$

where $\boldsymbol{\tau}^{(m)}$ denotes a $1 \times GV$ vector $(\tau^{(m)}_{1,1},\ldots,\tau^{(m)}_{G,V})^T$ (the superscript T denotes matrix transposition); $\text{Mul}(\cdot|\boldsymbol{\pi})$ is a multinomial distribution of one sample size with probability parameter $\boldsymbol{\pi}$; $\text{Beta}(\cdot|a,b)$ is a Beta distribution with prior sample size (a,b); $\boldsymbol{Y}^{(m)}_{j\cdot\cdot}$ is a $1 \times GV$ vector $(Y^{(m)}_{j,1,1},\ldots,Y^{(m)}_{j,V,G})^T$. Note that we truncate the number of views with sufficient large V and the number of feature clusters with G [2]. When $Y^{(m)}_{j,v,g} = 1$, feature j belongs to feature cluster g at view v. By default, we set the concentration parameters α_1 and α_2 to one.

Generative model for object clusters $\boldsymbol{Z}$

A subject cluster membership vector of object i in view v, denoted as $\boldsymbol{Z}_{i,v\cdot}$, is generated by

$$
\begin{aligned}
u_{k,v} &\sim \text{Beta}(\cdot|1,\beta),\ v = 1,2,\ldots,\ k = 1,2,\ldots \\
\eta_{k,v} &= u_{k,v} \prod_{t=1}^{k-1}(1-u_{t,v}), \\
\boldsymbol{Z}_{i,v\cdot} &\sim \text{Mul}(\cdot|\boldsymbol{\eta}_v),
\end{aligned}
$$

where $\boldsymbol{Z}_{i,v\cdot}$ is a $1 \times K$ (we take K sufficiently large) vector given by $\boldsymbol{Z}_{i,v\cdot} = (Z_{i,v,1}, \ldots, Z_{i,v,K})^T$. We set the concentration parameter β to one.

Our multiple clustering model is summarized in a graphical model of Figure 2. It clarifies causal links among relevant parameters and a data matrix.

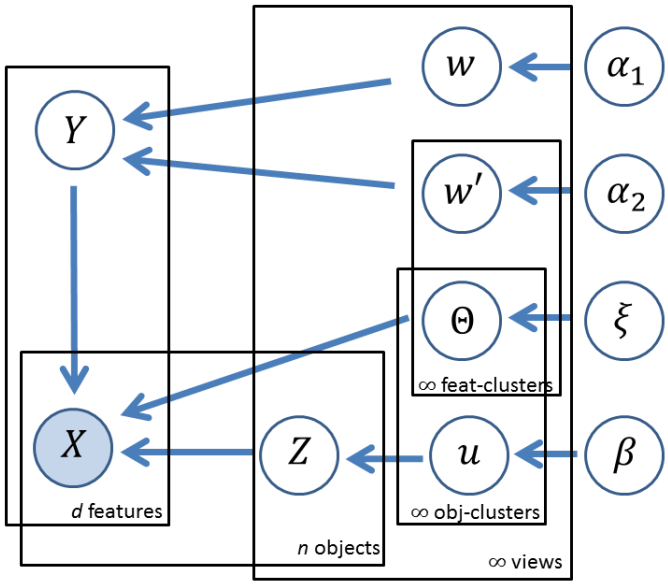

Figure 2: *Graphical model of relevant parameters in our multiple co-clustering model. Feat- and obj-cluster denotes feature and object cluster, respectively. Note that ξ denotes all hyperparameters for distributions of parameters $\boldsymbol{\Theta}$.*

2.2 Likelihood and prior distribution

We assume that each instance $X_{i,j}^{(m)}$ independently follows a certain distribution, conditional on $\boldsymbol{Y}$ and $\boldsymbol{Z}$. We denote $\boldsymbol{\theta}_{v,g,k}^{(m)}$ as parameters of distribution family m in the cluster block of view v, feature cluster g and object cluster k. Further denoting $\boldsymbol{\Theta} = \{\boldsymbol{\theta}_{v,g,k}^{(m)}\}_{v,g,k,m}$, the logarithm of complete-data likelihood of $\boldsymbol{X}$ is given by

$$\log p(\boldsymbol{X}|\boldsymbol{Y}, \boldsymbol{Z}, \boldsymbol{\Theta}) = \sum_{m,v,g,k,j,i} \mathbb{I}(Y_{j,v,g}^{(m)} = 1)\mathbb{I}(Z_{i,v,k} = 1) \log p(X_{i,j}^{(m)}|\theta_{v,g,k}^{(m)}),$$

where $\mathbb{I}(x)$ is an indicator function, i.e, returning 1 if x is true, and 0 otherwise. Note that the complete-data likelihood is not directly associated with $\boldsymbol{w} = \{w_v\}_v$, $\boldsymbol{w}' = \{w'^{(m)}_{g,v}\}_{g,v}$ and $\boldsymbol{u} = \{u_{k,v}\}_{k,v}$. The joint prior distribution for parameters is given by

$$p(\boldsymbol{w})p(\boldsymbol{w}')p(\boldsymbol{Y}|\boldsymbol{w}, \boldsymbol{w}')p(\boldsymbol{u})p(\boldsymbol{Z}|\boldsymbol{u})p(\boldsymbol{\Theta}).$$

2.3 Variational Inference

As in [8], we use variational Bayes EM for MAP (maximum a posteriori) estimation of $\boldsymbol{Y}$ and $\boldsymbol{Z}$. The logarithm of the marginal likelihood $p(\boldsymbol{X})$ is approximated using Jensen's inequality [12]:

$$\log p(\boldsymbol{X}) \;\geq\; \int q(\boldsymbol{\phi}) \log \frac{p(\boldsymbol{X}, \boldsymbol{\phi})}{q(\boldsymbol{\phi})} d\boldsymbol{\phi} = \mathcal{L}(q(\boldsymbol{\phi})), \tag{1}$$

where $q(\boldsymbol{\phi})$ is an arbitrary distribution for parameters $\boldsymbol{\phi}$. It can be shown that the difference between the left and right sides is given by the Kullback-Leibler divergence between $q(\boldsymbol{\phi})$ and $p(\boldsymbol{\phi}|\boldsymbol{X})$, i.e., $\mathbb{KL}(q(\boldsymbol{\phi}), p(\boldsymbol{\phi}|\boldsymbol{X}))$. Hence, our approach of choosing $q(\boldsymbol{\phi})$ is to minimize $\mathbb{KL}(q(\boldsymbol{\phi}), p(\boldsymbol{\phi}|\boldsymbol{X}))$, which is tractable to evaluate. In our model, we choose $q(\boldsymbol{\phi})$ that is factorized over different parameters (mean field approximation):

$$q(\boldsymbol{\phi}) \;=\; q_{\boldsymbol{w}}(\boldsymbol{w}) q_{\boldsymbol{w}'}(\boldsymbol{w}') q_{\boldsymbol{Y}}(\boldsymbol{Y}) q_{\boldsymbol{u}}(\boldsymbol{u}) q_{\boldsymbol{Z}}(\boldsymbol{Z}) q_{\boldsymbol{\Theta}}(\boldsymbol{\Theta}),$$

where each $q(\cdot)$ is further factorized over subsets of parameters, w_v, $w'^{(m)}_{g,v}$, $\boldsymbol{Y}^{(m)}_{j\cdot\cdot}$, $u_{k,v}$, $\boldsymbol{Z}_{i,v\cdot}$ and $\boldsymbol{\theta}^{(m)}_{v,g,k}$.

In general, the distribution $q_i(\phi_i)$ that minimizes $\mathbb{KL}(\prod_{l=1}^{L} q_l(\phi_l), p(\boldsymbol{\phi}|\boldsymbol{X}))$ is given by

$$q_i(\phi_i) \propto \exp\{\mathbb{E}_{-q_i(\phi)} \log p(\boldsymbol{X}, \boldsymbol{\phi})\},$$

where $\mathbb{E}_{-q_i(\phi)}$ denotes averaging with respect to $\prod_{l \neq i} q_l(\phi_l)$ [14]. Applying this property to our model, it can be shown that

$$\begin{aligned}
q_{\boldsymbol{w}}(\boldsymbol{w}) &= \prod_{v=1}^{V} \text{Beta}(w_v|\gamma_{v,1}, \gamma_{v,2}) \\
q_{\boldsymbol{w}'}(\boldsymbol{w}') &= \prod_{m=1}^{M} \prod_{v=1}^{V} \prod_{g=1}^{G} \text{Beta}(w^{(m)}_{g,v}|\gamma^{(m)}_{g,v,1}, \gamma^{(m)}_{g,v,2}) \\
q_{\boldsymbol{Y}}(\boldsymbol{Y}) &= \prod_{m=1}^{M} \prod_{j=1}^{d^{(m)}} \text{Mul}(\boldsymbol{Y}^{(m)}_{j\cdot\cdot}|\boldsymbol{\tau}^{(m)}_{j}) \\
q_{\boldsymbol{u}}(\boldsymbol{u}) &= \prod_{v=1}^{V} \prod_{k=1}^{K} \text{Beta}(u_{g,v}|\gamma_{k,v,1}, \gamma_{k,v,2}) \\
q_{\boldsymbol{Z}}(\boldsymbol{Z}) &= \prod_{v=1}^{V} \prod_{i=1}^{n} \text{Mul}(\boldsymbol{Z}_{i,v\cdot}|\boldsymbol{\eta}_{i,v}) \\
\log q_{\boldsymbol{\Theta}}(\boldsymbol{\Theta}) &= \sum_{m,v,g,k,j,i} \tau^{(m)}_{j,v,g} \eta_{i,v,k} \log p(X^{(m)}_{i,j}|\boldsymbol{\theta}^{(m)}_{v,g,k}) + \sum_{m,v,g,k} \log p(\boldsymbol{\theta}^{(m)}_{v,g,k}) + \text{constant},
\end{aligned}$$

where the hyperparameters except for $q_{\boldsymbol{\Theta}}(\boldsymbol{\Theta})$ are given by

$$\gamma_{v,1} \;=\; 1 + \sum_{m=1}^{M} \sum_{g=1}^{G} \sum_{j=1}^{d^{(m)}} \tau^{(m)}_{j,g,v}$$

$$
\begin{aligned}
\gamma_{v,2} &= \alpha_1 + \sum_{m=1}^{M}\sum_{t=v+1}^{V}\sum_{g=1}^{G}\sum_{j=1}^{d^{(m)}}\tau_{j,g,t}^{(m)} \\
\gamma_{g,v,1}^{(m)} &= 1 + \sum_{j=1}^{d^{(m)}}\tau_{j,g,v}^{(m)} \\
\gamma_{g,v,2}^{(m)} &= \alpha_2 + \sum_{t=g+1}^{G}\sum_{j=1}^{d^{(m)}}\tau_{j,t,v}^{(m)} \\
\gamma_{k,v,1} &= 1 + \sum_{i=1}^{n}\eta_{i,v,k} \\
\gamma_{k,v,2} &= \beta + \sum_{t=k+1}^{K}\sum_{i=1}^{n}\eta_{i,v,t} \\
\log\tau_{j,g,v}^{(m)} &= \sum_{k=1}^{K}\sum_{i=1}^{n}\eta_{i,v,k}\mathbb{E}_{q(\boldsymbol{\theta})}\Big[\log p(X_{i,j}^{(m)}|\boldsymbol{\theta}_{v,g,k}^{(m)})\Big] \\
&\quad +\psi(\gamma_{v,1}) - \psi(\gamma_{v,1}+\gamma_{v,2}) \\
&\quad +\sum_{t=1}^{v-1}\{\psi(\gamma_{t,2}) - \psi(\gamma_{t,1}+\gamma_{t,2})\} \\
&\quad +\psi(\gamma_{g,v,1}^{(m)}) - \psi(\gamma_{g,v,1}^{(m)}+\gamma_{g,v,2}^{(m)}) \\
&\quad +\sum_{t=1}^{G-1}\{\psi(\gamma_{t,v,2}^{(m)}) - \psi(\gamma_{t,v,1}^{(m)}+\gamma_{t,v,2}^{(m)})\} \\
&\quad +\text{constant} \\
\log\eta_{i,v,k} &= \sum_{m=1}^{M}\sum_{g=1}^{G}\sum_{j=1}^{d^{(m)}}\tau_{j,g,v}^{(m)}\mathbb{E}_{q(\boldsymbol{\theta})}\Big[\log p(X_{i,j}^{(m)}|\boldsymbol{\theta}_{v,g,k}^{(m)})\Big] \\
&\quad +\psi(\gamma_{k,v,1}) - \psi(\gamma_{k,v,1}+\gamma_{k,v,2}) \\
&\quad +\sum_{t=1}^{K-1}\{\psi(\gamma_{t,v,2}) - \psi(\gamma_{t,v,1}+\gamma_{t,v,2})\} \\
&\quad +\text{constant}, \qquad (2)
\end{aligned}
$$

where $\mathbb{E}_{q(\boldsymbol{\theta})}$ denotes averaging with respect to the corresponding $q(\boldsymbol{\theta})$ of $\boldsymbol{\theta}_{v,g,k}^{(m)}$; $\psi(\cdot)$ denotes the digamma function defined as the first derivative of logarithm of gamma function. Note that $\tau_{j,g,v}^{(m)}$ is normalized over pairs (g,v) for each pair (j,m), while $\eta_{i,v,k}$ normalized over k for each pair of (i,v). Observation models and priors of parameters $\boldsymbol{\Theta}$ are specified in the following section.

Algorithm 1 Variational Bayes EM for multiple co-clustering

Input: data matrices $\boldsymbol{X}^{(1)}, \ldots, \boldsymbol{X}^{(M)}$.
for $s = 1$ **to** S **do**
  Randomly initialize $\{\boldsymbol{\tau}^{(m)}\}_m$ and $\{\boldsymbol{\eta}_v\}_v$.
  repeat
    -Update the hyperparameters of relevant distribution families for $q_{\boldsymbol{\Theta}}(\boldsymbol{\Theta})$.
    -Update the hyperparameters for $q_{\boldsymbol{w}}(\boldsymbol{w})$, $q_{\boldsymbol{w}'}(\boldsymbol{w}')$, $q_{\boldsymbol{Y}}(\boldsymbol{Y})$, $q_{\boldsymbol{u}}(\boldsymbol{u})$, and $q_{\boldsymbol{Z}}(\boldsymbol{Z})$.
  until L in Eqs.(3) converges.
  Keep $L(s) = L$
end for
$s^* = \text{argmax}_s L(s)$
Output: MAP for $\boldsymbol{Y}$ and $\boldsymbol{Z}$ in the run s^*.

2.4 Observation models

For observation models, we consider Gaussian, Poisson, and Categorical/multinomial distributions. For each cluster block, we fit a univariate distribution of these families with the assumption that features within the cluster block are independent. We assume conjugate priors for the parameters of these distribution families. Variational inference and updating equations are basically the same as in [8] (See Appendix A).

2.5 Algorithm

With the updating equations of the hyperparameters, the variational Bayes EM proceeds as follows. First, we randomly initialize $\{\boldsymbol{\tau}^{(m)}\}_m$ and $\{\boldsymbol{\eta}_v\}_v$, and then alternatively update the hyperparameters until the lower bound $\mathcal{L}(q(\boldsymbol{\phi}))$ in Eqs.(1) converges. This yields possible optimal distributions $q_{\boldsymbol{Y}}(\boldsymbol{Y})$ and $q_{\boldsymbol{Z}}(\boldsymbol{Z})$, which immediately gives MAP estimates for $\boldsymbol{Y}$ and $\boldsymbol{Z}$. We repeat this procedure a number of times, and choose the best solution with the largest lower bound. The algorithm is outlined in Algorithm 1. Note that the lower bound $\mathcal{L}(q(\boldsymbol{\phi}))$ is given by

$$\mathcal{L}(q(\boldsymbol{\phi})) = \int q(\boldsymbol{\phi}) \log p(X|\boldsymbol{\phi}) d\boldsymbol{\phi} - \mathbb{KL}(q(\boldsymbol{\phi}), p(\phi)), \tag{3}$$

where both terms on the right side can be derived in closed form. It can be shown that this monotonically increases as $q(\boldsymbol{\phi})$ is optimized.

2.6 Time complexity

For simplicity, we consider time complexity of our algorithm for a single run. If we assume that the number of required iterations for convergence is the same, the time complexity of the algorithm is equivalent to the number of operations for updating the relevant parameters. In that case, as can be seen in the updating equations in Eqs.(2) and Appendix A, the time complexity is just $O(nd)$ where n and d are the number of objects and the number of features (we fix the number

of views and clusters). This enhances efficiency in applying our multiple co-clustering method to high-dimensional data. We return to this point in Section 4.4 to compare other multiple clustering methods.

2.7 Model representation

Our multiple co-clustering model is sufficiently flexible to represent different clustering models because the number of views and the number of feature-/object clusters are derived in a data-driven approach. For instance, when the number of views is one, the model coincides with a co-clustering model; when the number of feature clusters is one for all views, it matches the restricted multiple clustering model. Furthermore, when the number of views is one and the number of feature clusters is the same as the number of features, it matches conventional mixture models with independent features. Moreover, our model can detect non-informative features that do not discriminate between object clusters. In such a case, the model yields a view in which the number of object clusters is one. The advantage of our model is to automatically detect such underlying data structures.

2.8 Missing values

Our multiple co-clustering model can easily handle missing values. Suppose that the missing entries occur at random, which may depend on the observed data, but not the missing ones (i.e., MAR, missing at random). We can deal with such missing values in a conventional Bayesian way, in which missing entries are considered as stochastic parameters [7]. In our model, this procedure is simply reduced to ignoring these missing entries when we update the hyperparameters. This is because (univariate) instances within a cluster block are assumed to be independent; hence the log-likelihood in Eqs.(1) is given by

$$\log p(\boldsymbol{X}^{obs}|\boldsymbol{Y},\boldsymbol{Z},\boldsymbol{\Theta}) = \sum_{m,v,g,k,j,i} \mathbb{I}(Y_{j,v,g}^{(m)}=1)\mathbb{I}(Z_{i,v,k}=1)\mathbb{I}((i,j)^{(m)} \in o) \log p(X_{i,j}^{(m)}|\theta_{v,g,k}^{(m)}),$$

where $\mathbb{I}((i,j)^{(m)} \in o)$ is an indicator for the status of availability of the data cell of object i and feature j for distribution family m (1 when it is available, and 0 otherwise); $\boldsymbol{X}^{obs}$ a subset of $\boldsymbol{X}$ that consists of the observed data only.

3 Simulation study on synthetic data

In this section, we report on a simulation study to evaluate the performance of our method. To the best of our knowledge, there is no algorithm in the literature that allows mixing of different types of features, as we have so far modeled. Hence, we compare the performance of our multiple co-clustering method only with co-clustering and restricted multiple clustering methods, which we model to accommodate different types of features. We set the hyperparameters α_1, α_2, and β relevant for generating views, feature clusters, and sample clusters in Section 2.1 to one, and the hyperparameters relevant for observations models to those specified in Appendix A. Note that we use this setting for further application of our method in Section 4. For data generation, we

Table 1: Summary of results of simulation study on synthetic data: Recovery of true object cluster structure and views evaluated in terms of mean values of adjusted Rand Index.

		Object clustering			Views		
Factors		**Mul**	Co	rMul	**Mul**	Co	rMul
Number of objects	20	**0.30** (0.01)	0.21 (0.00)	0.13 (0.01)	0.05 (0.01)	**0.27** (0.00)	0.11 (0.00)
	50	**0.81** (0.01)	0.23 (0.00)	0.51 (0.01)	**0.75** (0.01)	0.48 (0.00)	0.37 (0.01)
	100	**0.84** (0.01)	0.21 (0.00)	0.69 (0.01)	**0.77** (0.01)	0.54 (0.00)	0.43 (0.01)
Number of features	10	**0.46** (0.01)	0.21 (0.00)	0.26 (0.01)	0.29 (0.01)	**0.33** (0.01)	0.22 (0.01)
	50	**0.75** (0.01)	0.22 (0.00)	0.49 (0.01)	**0.66** (0.01)	0.47 (0.00)	0.33 (0.01)
	100	**0.73** (0.01)	0.21 (0.00)	0.58 (0.01)	**0.62** (0.01)	0.49 (0.00)	0.37 (0.01)
Ratio of missings	0	**0.70** (0.01)	0.22 (0.00)	0.49 (0.01)	**0.58** (0.01)	0.45 (0.00)	0.33 (0.01)
	0.1	**0.65** (0.01)	0.22 (0.00)	0.44 (0.01)	**0.53** (0.01)	0.43 (0.01)	0.30 (0.01)
	0.2	**0.59** (0.01)	0.21 (0.00)	0.40 (0.01)	**0.48** (0.01)	0.41 (0.01)	0.28 (0.01)

(a) Mul, Co, and rMul denote our multiple co-clustering, co-clustering and restricted multiple clustering methods.
(b) Digits without parenthesis denotes mean values of adjusted Rand Index over 9×100 = 900 datasets for a corresponding factor.
(c) Digits in parenthesis denote standard errors of the mean.
(d) The digits for the best performance among the three methods are in bold.

fixed the number of views to three and the number of object clusters to two, three, and four in views 1-3, respectively. The number of feature clusters was set to two in all views (Figure 3a). We manipulated the number of features (per view *and* distribution family) (10, 50, 100), the number of objects (20, 50, 100), and the proportion of (uniformly randomly generated) missing entries (0, 0.1, 0.2). We included three types of mixtures of distributions: Gaussian, Poisson, and Categorical. Memberships of views were evenly assigned to features for each distribution family, and the feature and object cluster memberships were uniformly randomly allocated. The distribution parameters for each cluster block were fixed as in the legend of Figure 3. We generated 100 datasets for each setting, which resulted in 100×27=2700 datasets.

We evaluated the performance of recovering the true cluster structure by means of an adjusted Rand index (ARI) [10]: When ARI is one, recovery of the true cluster structure is perfect. When ARI is close to zero, recovery is almost random. Specifically, we focused on recovery of mem-

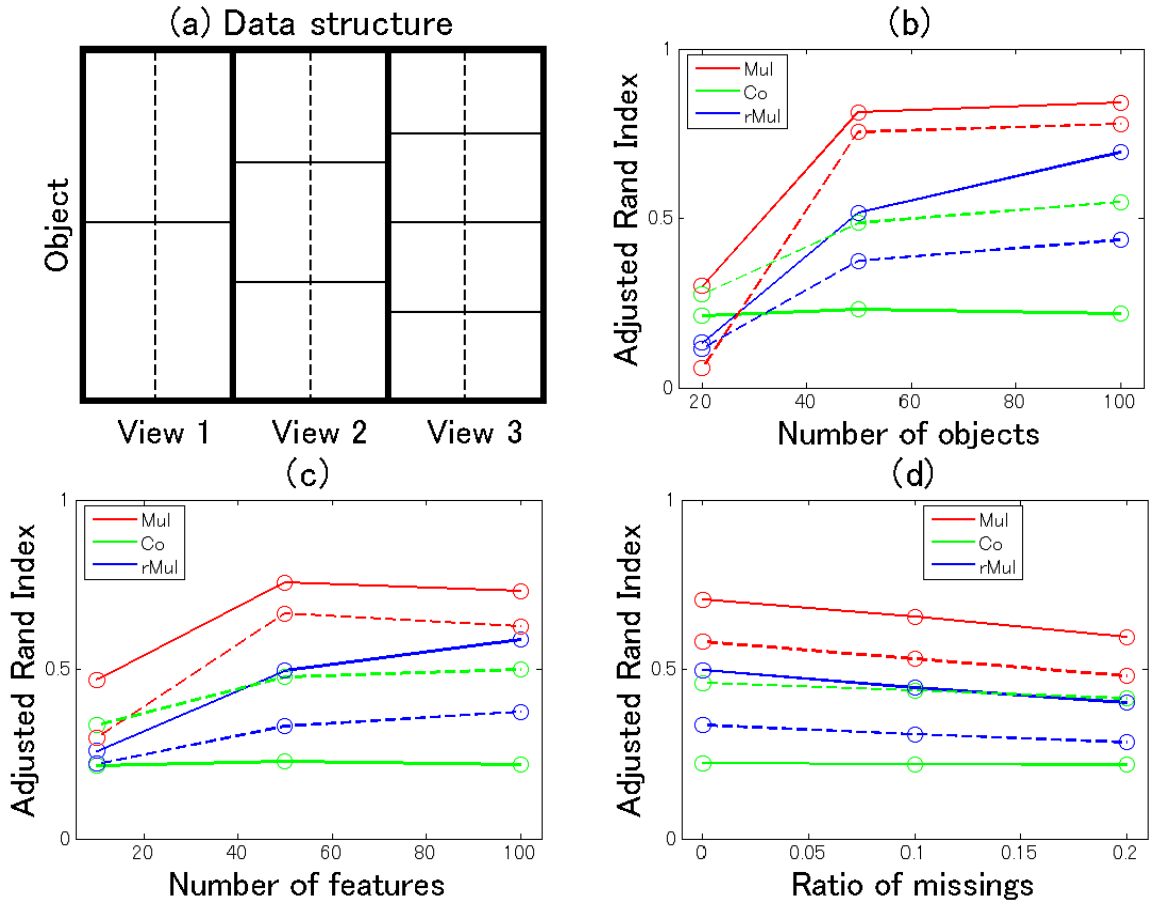

Figure 3: *Panel (a): Data structure for the simulation study. Each view has two feature clusters (separated by a dashed line) for each type of features of Gaussian, Poisson and Categorical. For Gaussian, means are set to (0, 4 ; 1, 3) ((0, 4) for top left and right cluster blocks, (1, 3) for bottom left and right cluster blocks) for view 1; (0, 5 ; 1, 4 ; 2, 3) for view 2; (0, 6 ; 1, 5 ; 2, 4 ; 3, 3) for view 3. The standard deviation is fixed to one. Similarly, for Poisson, the parameter λ is set to (1, 2 ; 2, 1), (1, 3 ; 2, 2 ; 3, 1), (1, 4 ; 2, 3 ; 3, 2 ; 4,1). For categorical (binary), probability for success is (0.1, 0.9 ; 0.1, 0.9),(0.1, 0.9 ; 0.5, 0.5; 0.9, 0.1), (0.1, 0.9 ; 0.4, 0.6 ; 0.6, 0.4 ; 0.9, 0.1). Panels (b)-(d): Performance of the multiple co-clustering method ('Mul', red), the co-clustering method ('Co', green), and the restricted multiple clustering method ('rMul', blue). Solid lines are for recovery of object cluster solutions, while dashed lines are for recovery of views. The results are summarized with respect to the number of objects (b), the number of features (c) and the proportions of missing entries (d).*

berships of views, and memberships of object clusters. Since the numbering of views is arbitrary, it is not straightforward to evaluate recovery of the true object cluster solutions (the correspondence between the yielded and the true object cluster solutions is not clear). Hence, to evaluate the performance of object cluster solutions, we first evaluated ARIs for all combinations of the true object clusters and yielded object cluster solutions, and then found the maximum ARI for each true object cluster solution. Lastly, we averaged the ARIs over views. In this manner, we evaluated the performance for the multiple co-clustering method and the restricted multiple clustering. The

co-clustering method yields only a single object cluster solution; hence we averaged ARIs between the true object cluster solutions and this solution.

The performance of the multiple co-clustering method is reasonably good: performance of the recovery of views (red dashed line) and object clusters (red solid line) solutions improves as the number of objects increases (Table 1 and Figure 3b). Regarding the number of features, the performance improves as the number of features increases from 20 to 50, but there is no improvement from 50 to 100. This is possibly because in our simulation setting, each feature does not clearly discriminate between object clusters; hence, adding more features does not necessarily improve the recovery of views (hence, the recovery of object cluster solutions). Lastly, when the ratio of missing entries increases, the performance just becomes slightly worse, which suggests that our method is relatively robust to missing entries.

As a whole, the multiple co-clustering method outperforms the co-clustering and the restricted multiple clustering methods. The performance of the co-clustering method is poor because it does not fit the multiple clustering structure. On the other hand, the restricted multiple clustering method can potentially fit each object cluster structure; hence, it performs somewhat well in this regard (but, not for recovery of the true memberships of views).

4 Application to real data

To test our multiple co-clustering method on real data, we consider three datasets: facial image data, cardiac arrhythmia data, and depression data. For facial image and cardiac arrhythmia data, the (possible) true sample clustering label is available, which enables us to evaluate clustering performance of our multiple co-clustering method. We compare the performance with the restricted multiple clustering method and two state-of-the-art multiple clustering methods: the constrained orthogonal average link algorithm (COALA, [1]) and the decorrelated K-means algorithm [11]. These state-of-the-art methods aim to detect dissimilar multiple sample clustering solutions without partitioning of features. COALA is based on a hierarchical clustering algorithm, while decorrelated K-means is based on a K-means algorithm. The two methods also differ in the way to detect views: COALA iteratively identifies views while decorrelated K-means simultaneously does so. For both methods, we need to set the number of views and the number of sample clusters. In this experiment, we set these to the (possible) true numbers. For the depression data, no information is available on true cluster structure. Hence, we focus mainly on implications drawn from the data by our multiple co-clustering method, rather than evaluating the performance of recovery of true cluster structure.

4.1 Facial image data

The first dataset contains facial image data from the UCI KDD repository (`http://archive.ics.uci.edu/ml/datasets.html`), which consists of black and white images of 20 different persons with varying configurations (Figure 4): *eyes* (wearing sunglasses or not), *pose* (straight, left, right, up), and *expression* (neutral, happy, sad, angry). This dataset served as a benchmark for multiple clustering in several papers [8, 1]. Here, we focus on the quarter-resolution images (32×30) of this dataset, which results in 960 features. For simplicity, we consider two subsets

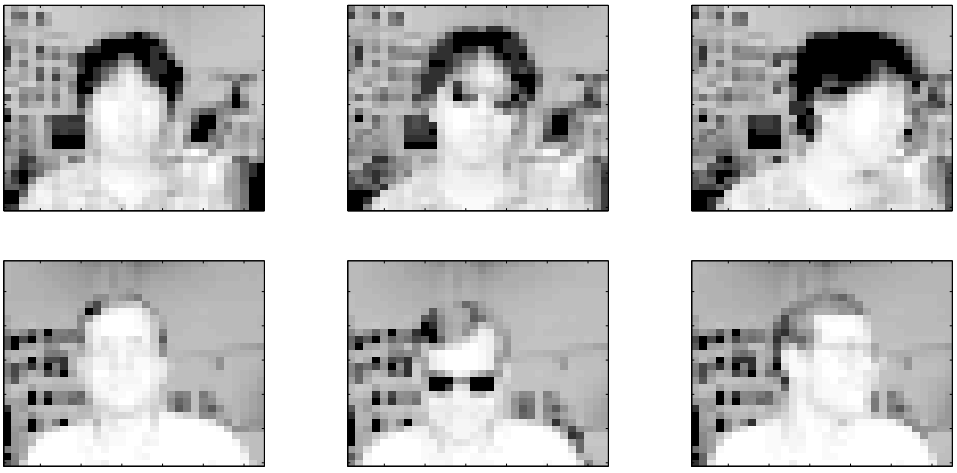

Figure 4: *Samples from the facial image data: The first row represents person 'an2i' with configurations of (no sunglasses, straight pose and neutral expression), (sunglasses, straight pose, angry expression) and (no sunglasses, left pose, happy expression) from left to right columns. The second row for person 'at33' with the same patterns of configuration.*

of these images: data 1 consisting of a single person (named 'an2i') with varying *eyes*, *pose* and *expression* (data size: 32×960); data 2 consisting of two persons (in addition to 'an2i', we include person 'at33', data size: 64×960). We use these datasets without pre-processing.

The facial image data has multiple clustering structures that can be characterized by all of the features (global) or some of the features (local). Identification of persons (hereafter, *useid*) may be related to global information of the image (all features), while *eyes*, *pose* and *expression* are based on local information (a subset of features). Here, we focus only on *pose*, which is a relatively easy aspect to detect [15]. Since COALA and decorrelated K-means methods do not explicitly model relevant features for sample clustering, they can potentially capture a multiple clustering structure based on either global or local information of such a dataset. On the other hand, our multiple co-clustering model is based on a partition of features, which implicitly assumes that a possible sample clustering structure is based on non-overlapping local information. Our interest in this application is to examine the performance of our multiple clustering method using such data.

To evaluate performance, we focus only on sample clustering solutions. We base our evaluation criterion on recovery of structures of *useid* and *pose* (*useid* is applicable only for Data 2), which is measured by the maximum value of an adjusted Rand Index between the true sample structure in question and resulting sample clustering solutions. We discuss the results for each data in the following sub-sections.

4.1.1 Results: Data 1

Our multiple co-clustering method yielded nine sample clusterings (i.e., nine views), one of which is closely related to *pose* with an adjusted Rand Index of 0.84 (Figure 5a, p <0.001 by permutation test, and Table 2 for the contingency table between true clusters and resultant sample clusters).

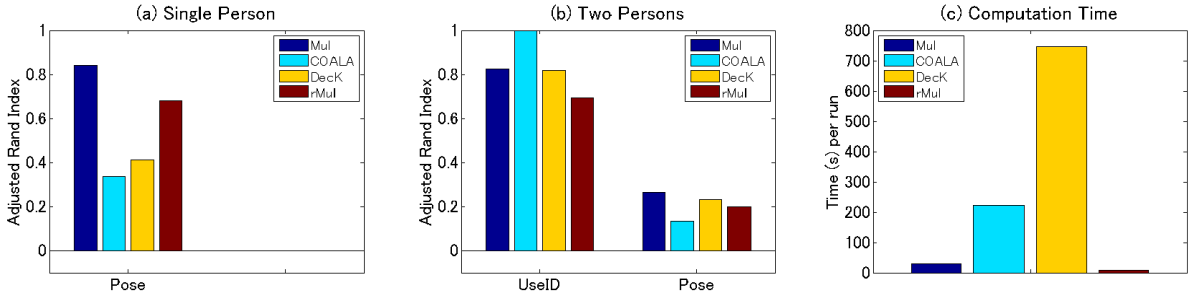

Figure 5: *Performance on sample clusterings for the facial image data. Panel (a) for the subset (Data 1) of a single person ('an2i'). Performance on pose for four clustering methods, i.e., multiple co-clustering (Mul), COALA, decorrelated K-means (DecK), and restricted multiple clustering (rMul) are evaluated by adjusted Rand Index of sample clustering solutions. Panel (b) for the subset (Data 2) of two persons ('an2i' and 'at33'). Performance is evaluated on useid and pose. Note that to match true and yielded views, we evaluated the maximum value of adjusted Rand index between the true sample clustering in question and the yielded sample cluster solutions. The number of initializations is 500 for the multiple co-clustering, decorrelated K-means and the restricted multiple clustering. Panel (c) for computation time (seconds) per single run of each clustering method.*

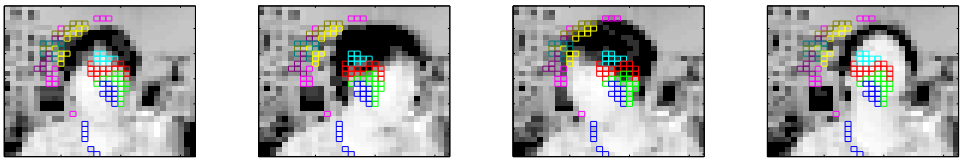

Figure 6: *Selected features by our multiple co-clustering method for person 'an2i' in the facial image data: Pixels surrounded by color boxes are the selected features that yielded the relevant sample clustering to pose. Color denotes a particular feature cluster.*

Our method outperforms COALA and decorrelated K-means methods (the performance of both methods is similar), and performs slightly better than the restricted multiple clustering method.

Next, we analyze features that are relevant to the sample clustering based upon *pose*. Note that our multiple co-clustering method yields information about features relevant to a particular sample clustering solution in an explicit manner while COALA and decorrelated K-means do not. Our method yielded the pixels (features) relevant to the cluster assignment, concentrating around subregions in the right part of head and the left part of face (Figure 6). This allows us to conclude that these subregions are very sensitive to different poses.

4.1.2 Results: Data 2

Our multiple co-clustering method yielded ten sample clusterings, three of which were closely related to *useid* (identification of person) and *pose* with adjusted Rand Indices of 0.82 (p <0.001) and

Table 2: *Results of sample clustering for data 1 of the facial image data: Contingency table of the true labels (pose) and yielded clusters of multiple co-clustering (Mul), COALA, decorrelated K-means (DecK), and restricted multiple (rMul) method from (a) to (d). T1, T2, T3 and T4 are true classes of pose (straight, left, right and up); C1, C2, C3, C4 and C5 are yielded clusters for each method.*

(a) Mul

	T1	**T2**	**T3**	**T4**
C1	8	0	0	2
C2	0	8	0	0
C3	0	0	8	0
C4	0	0	0	6

(b) COALA

	T1	**T2**	**T3**	**T4**
C1	2	0	6	1
C2	6	0	1	7
C3	0	5	1	0
C4	0	3	0	0

(c) DecK

	T1	**T2**	**T3**	**T4**
C1	0	0	2	2
C2	2	3	1	2
C3	4	4	4	3
C4	2	1	1	1

(d) rMul

	T1	**T2**	**T3**	**T4**
C1	0	8	0	0
C2	0	0	8	0
C3	1	0	0	5
C4	4	0	0	2
C5	3	0	0	1

Table 3: *Results for data 2 of the facial image data: Contingency table of the true labels (useid) and yielded clustering of the multiple co-clustering (Mul), COALA, decorrelated K-means (DecK), and restricted multiple (rMul) method from (a) to (d). T1 and T2 are true classifications (an2i, at33); C1, C2, C3 and C4 are yielded clusters.*

(a) Mul

	T1	**T2**
C1	32	0
C2	0	25
C3	0	7

(b) COALA

	T1	**T2**
C1	32	0
C2	0	32

(c) DecK

	T1	**T2**
C1	3	32
C2	29	0

(d) rMul

	T1	**T2**
C1	32	0
C2	0	15
C3	0	13
C4	0	4

0.26 (p <0.001), respectively (Figure 5b, and Tables 3 and 4 for the contingency tables for *useid* and *pose*, respectively). To compare with COALA, our multiple co-clustering method performed a bit poorly for detecting *useid*, while it performed better for *pose*. On the other hand, the performance of our method is comparable to the decorrelated K-means method. Further, it performed slightly better than the restricted multiple clustering method.

The most relevant pixels for *useid* concentrate near the right part of face, and the background (Figure 7). This can be interpreted to mean that the difference in hair style may be an important factor to distinguish between these two persons. In addition, an apparent difference in their rooms (background) also serves as a discriminating factor.

Table 4: *Results of sampling-clustering for data 2 of the facial image data: Contingency table of the true labels (pose) and yielded clustering of the multiple co-clustering (Mul), COALA, decorrelated K-means (DecK), and restricted multiple (rMul) method from (a) to (d). T1, T2, T3 and T4 are true classes of pose (straight, left, right and up); C1, ..., C7 are yielded results for each method.*

(a) Mul

	T1	**T2**	**T3**	**T4**
C1	7	8	0	8
C2	1	0	0	8
C3	0	1	7	0
C4	0	7	1	0
C5	7	0	0	0
C6	0	0	5	0
C7	1	0	3	0

(b) COALA

	T1	**T2**	**T3**	**T4**
C1	8	1	1	6
C2	0	7	0	0
C3	0	0	9	0
C4	8	8	6	10

(c) DecK

	T1	**T2**	**T3**	**T4**
C1	7	4	0	2
C2	0	3	14	3
C3	6	3	0	10
C4	3	6	2	1

(d) rMul

	T1	**T2**	**T3**	**T4**
C1	5	3	2	4
C2	5	0	8	0
C3	0	0	6	4
C4	1	1	0	8
C5	2	7	0	0
C6	3	5	0	0

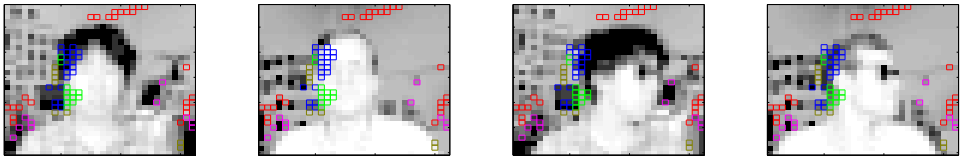

Figure 7: *Samples from image datasets for person 'an2i' and 'at33': Pixels surrounded by color boxes are selected features that yielded relevant sample clustering to useid in data2. Image configurations are ('an2i', non sunglass, straight), ('at33', non sunglass, straight), 'an2i', sunglass, left), and ('at33', sunglass, left), respectively. Expression is neutral for all samples. In these examples, the multiple clustering method correctly identified these persons.*

4.2 Cardiac Arrhythmia data

Next, we apply our multiple co-clustering method to Cardiac Arrhythmia data (UCI KDD repository). Unlike the facial image data in the previous section, this dataset does not necessarily have multiple sample clustering structures (indeed, such information is not available). However, the multiple co-clustering method should be able to automatically select relevant features.

The original dataset consisted of 452 samples (subjects) and 279 features that comprise various cardiac measurements and personal information such as sex, age, height, and weight (See more

detail in [9]). Some of these features are numerical (206 features) and others are categorical (73 features). Further, there are a number of missing entries in this dataset. Beside these features, a classification label is available, which classifies the subjects into one of 16 types of arrhythmia. For simplicity, we focus only on three types of arrhythmia of similar sample size: *Old Anterior Myocardial infarction* (sample size 15), *Old Inferior Myocardial Infarction* (15) and *Sinus Tachycardia* (13). The objective in this section is to examine recovery performance among these three types of arrhythmia.

Application of COALA and decorrelated K-means methods to this dataset is problematic because these methods do not allow for categorical features nor missing entries. Hence, we use the following heuristic procedure to pre-process the data: Re-code a binary categorical feature using a numerical feature (taking values 0 or 1); replace missing entries with mean values of features. Recall that these problems do not arise with our multiple co-clustering method.

Results

Our multiple co-clustering method yielded nine sample clusterings (i.e., nine views). The maximum adjusted Rand Index between the true labels and resultant clusters is 0.56 (p <0.001, Figure 8). On the other hand, the maximum Rand Index for COALA, decorated K-means and restricted multiple clustering methods are 0.02, 0.49 and 0.39, respectively.

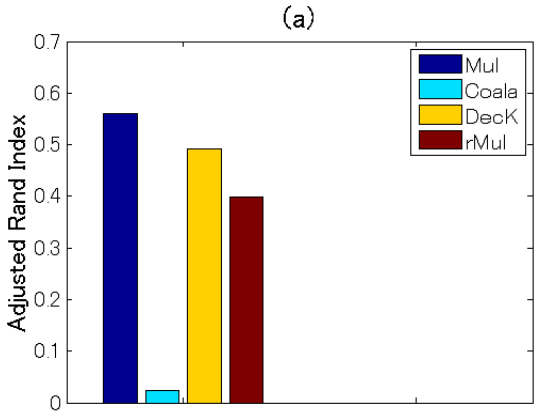

Figure 8: *Results of multiple clustering for the cardiac arrhythmia data: Comparison of performance on subject clustering in terms of adjusted Rand Index among multiple co-clustering, COALA, decorrelated K-means and restricted multiple clustering methods*

Further, we examine subject clustering results more in detail. For our multiple co-clustering method, the subject clusters C1, C2, and C3 distinguish the three symptoms well (corresponding to T2, T1, and T3, respectively, Table 5). On the other hand, such a distinction is totally or partially ambiguous for the other methods. In the case of COALA, clustering results seem to be degenerate, yielding two tiny clusters (C2 and C3). For decorrelated K-means, the distinction among T1, T2, and T3 is partially ambiguous. There is a clear correspondence between C1 and T1, but C2 is a

Table 5: *Results of sample clustering for the cardiac arrhythmia data: Contingency table of the true labels and yielded clustering of the multiple co-clustering (Mul), decorrelated K-means (DecK), and restricted multiple (rMul) method from (a) to (d). T1, T2, and T3 are true classes of arrhythmia (Old Anterior Myocardial Infarction, Old Inferior Myocardial infarction, and Sinus Bradycardy, respectively); C1, C2, C3 and C4 are yielded results for each method.*

(a) Mul

	T1	**T2**	**T3**
C1	0	14	6
C2	14	0	0
C3	1	1	5
C4	0	0	2

(b) COALA

	T1	**T2**	**T3**
C1	15	15	10
C2	0	0	2
C3	0	0	1

(c) DecK

	T1	**T2**	**T3**
C1	14	0	0
C2	0	0	1
C3	1	15	12

(d) rMul

	T1	**T2**	**T3**
C1	14	0	2
C2	1	5	4
C3	0	3	6
C4	0	7	1

tiny cluster, and C3 does not distinguish between T2 and T3. A similar observation is made for the restricted multiple clustering method.

4.3 Depression data

Lastly, we apply our multiple co-clustering method to depression data, which consists of clinical questionnaires and bio-markers for healthy and depressive subjects. The objective here is to explore ways of analyzing the results from our multiple co-clustering method in a real situation where the true subject-cluster structures are unknown. The depression data comprise 125 subjects (66 healthy and 59 depressive) and 243 features (Table 6 in Appendix B) that were collected at a collaborating university. Among these features, there are 129 numerical (e.g., age, severity scores of psychiatric disorders) and 114 categorical features (e.g., sex, genotype) with a number of missing entries. Importantly, these data were collected from subjects in three different phases. The first phase was when depressive subjects visited a hospital for the first time. The second phase was 6 weeks after subjects started medical treatment. The third phase was 6 months after the onset of the treatment. For healthy subjects, relevant data for the second and the third phases were not available (dealt as missing entries in the data matrix). To distinguish between these phase differences, we denote features in the second and the third phases with endings of $6w$ and $6m$, respectively. Further, we did not include the label of health/depression status for clustering. We used it only for interpretation of results. We assumed that numerical features follow mixtures of Gaussian distributions in our model. To pre-process numerical features, we standardized each of them using means and standard deviations of available (i.e., nonmissing) entries.

Results

Our multiple co-clustering method yielded seven views. The majority of features are allocated to two views (view 1 and view 2 in Figure 9a). The number of subject clusters ranges from one to five (Figure 9b). We analyze these cluster results more in detail, focusing on view 1 and view 2. View 1 has two feature clusters for numerical features, in which the majority of features are related to DNA methylation of CpG sites of the trkb and htr2c genes with a number of missing entries (Figure 10a).

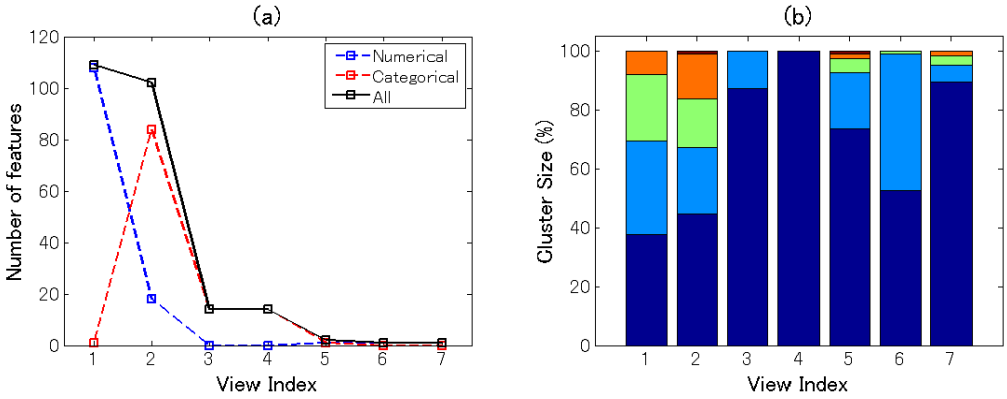

Figure 9: *Results of the multiple co-clustering method for clinical data of depression. Panel (a): Number of features (in black) in each view with numerical features in blue and categorical features in red. Panel (b): cluster size (percentage of subjects) for subject clusters in each view.*

For better visualization of this view, we remove methylation-related features (Figure 10b). Among these two (numerical) feature clusters, feature cluster 1 does not discriminate well between the yielded subject clusters (Figure11a), while feature cluster 2 does well (Figure 11b). Hence, subject clustering in this views is largely characterized by features in feature cluster 2 (BDI26w, BDI26m, PHQ96w, PHQ96m, HRSD176w, HRSD216w, CATS:total, CATS:N, and CATS:E). The first six features are related to psychiatric disorder scores at six weeks (features ending with -$6w$) and six months (features ending with -$6m$) after the onset of depression treatment. Hence, we can interpret this to reflect treatment effects. On the other hand, CATS:total, CATS:N, and CATS:E are related to abusive experiences in the subject's childhood. Hence, these features are available before the onset of treatment. These data attributes suggest that it is possible to predict treatment effect by using features related to child abuse experiences. In particular, the distribution pattern in subject cluster 3 (Figure 10b) is remarkably different from those in the remaining subject clusters (Figure 11b).

In view 2, healthy and depressive subjects are well separated. The first subject cluster is for healthy subjects. The second is intermediate, and the third and fourth are for depressive subjects (Figure 10c). Relevant numerical features are: JART, GAF, GAF6w, PANASP, PANASP6w, LES:total, LES:N, BAS, E, O, A, C, and Rep. in feature cluster 1 and BDI2, PHQ9, SHAPS, PANASN, and STAI in feature cluster 2. This result is quite reasonable, because the majority of these features are scores from clinical questionnaires that evaluate depressive disorders either negatively (smaller values in feature cluster 1) or positively (larger values in feature cluster 2).

4.4 Comparison of time complexity

Finally, we briefly discuss complexity of the clustering methods. Except for COALA, we need to run clustering methods (i.e., multiple co-clustering, decorrelated K-means, and restricted multiple clustering) with a number of random initializations for their parameters, and choose the optimal solution. Hence, computation time depends on the number of initializations. To compare complexity of computation, we make several assumptions. First, we focus on a single run of each method.

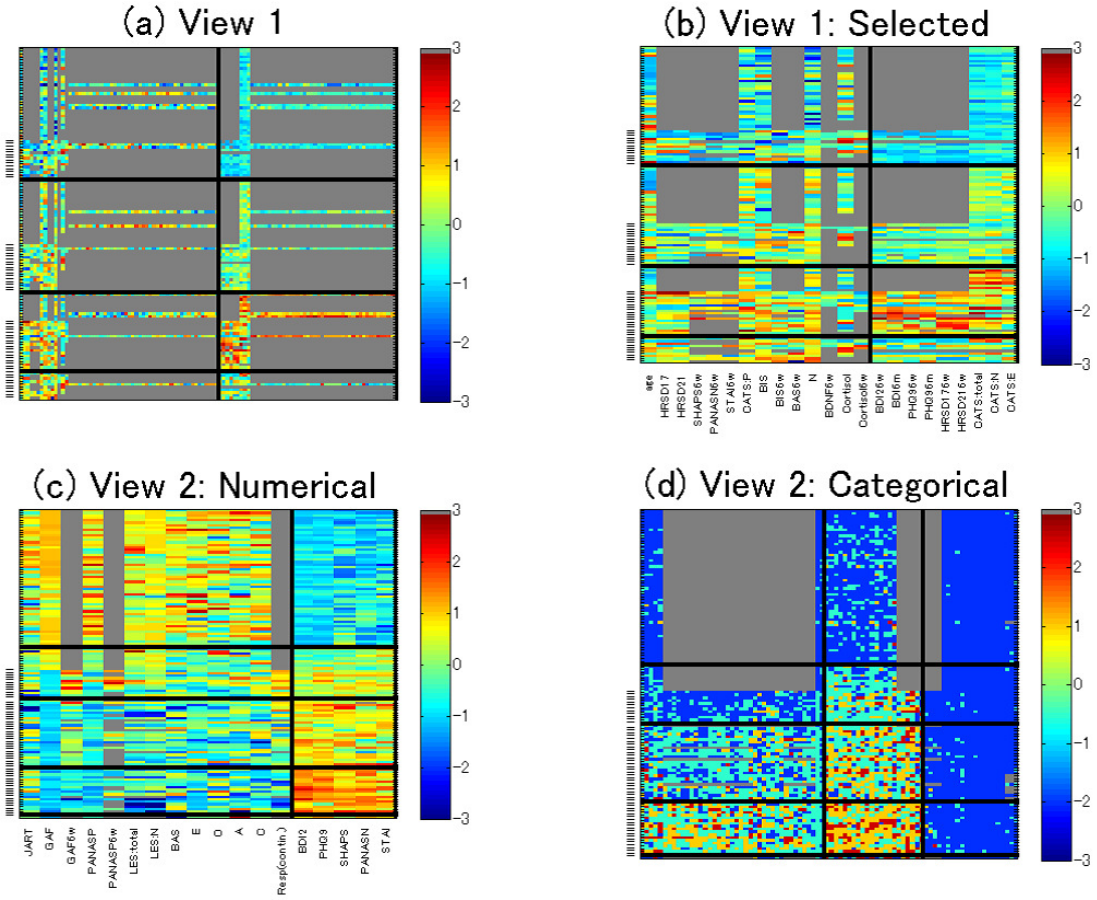

Figure 10: *Visualizations of views yielded by our multiple co-clustering method. Panels (a)-(b): Heatmaps of views 1. The x-axis denotes numerical features, and the y-axis denotes subjects. A depressive subject is indicated by a hyphen in left. The subject clusters are sorted in the order of cluster size. Panel (b) is a copy of panel (a) after removing methylation related features (those having a large number of missing entries). Panels (c)-(d): Heatmaps of views 2. Panel (c) contains numerical features while panel (d) contains categorical ones. The subject clusters are sorted in the descending order of the proportion of depressive subjects. For these panels, the subjects within a subject cluster are sorted in the order of healthy and depressive subjects. On the other hand, feature clusters are sorted in the order of feature clusters in the order of feature cluster size. Note that for categorical features the color is arbitrary and that missing entries are in gray.*

Second, we assume the same number of iterations for convergence. Third, we assume that the numbers of views and clusters are fixed. In such a setting, time complexity for our multiple clustering method (as well as the restricted multiple clustering) is $O(nd)$, where n and d are the number of samples and the number of features, respectively (Section 2.6). This suggests that the complexity is just linear when either n or d is fixed. On the other hand, the complexities of COALA and decorrelated K-means are $O(n^2 \log n + n^2 d)$ and $O(nd + d^3)$, respectively, based on their typical algorithms [1, 11]. These results imply that the complexity of our multiple co-clustering is generally less than

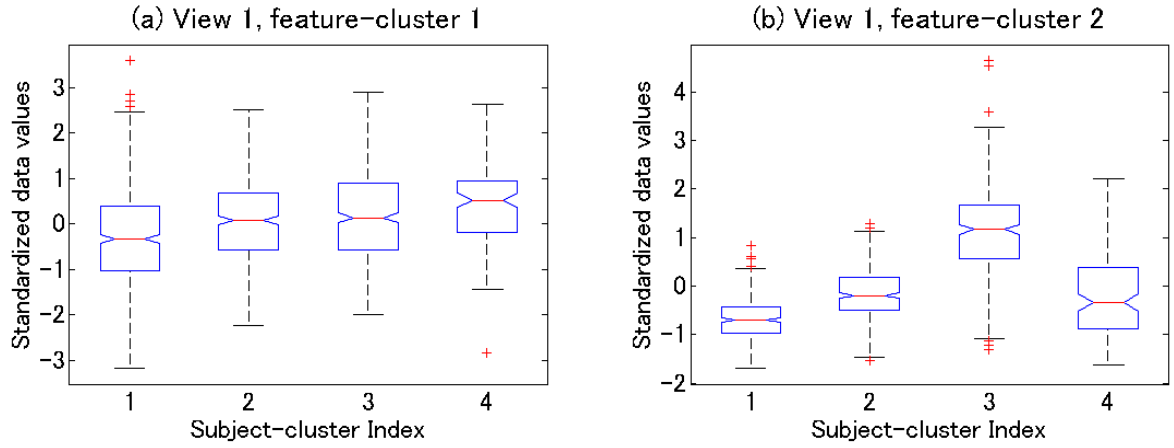

Figure 11: *Distributions of the standardized data within feature clusters 1 (Panel (a)) and 2 (Panel (b)). X-axis denotes subject cluster index. All relevant entries except for missing ones are accommodated in each box.*

those of COALA and decorrelated K-means, suggesting superior efficiency of the present method. Indeed, in the simulation of facial image data, our multiple clustering method requires less time per run than COALA and decorrelated K-means (Figure 5c).

5 Discussion

We propose a novel method of multiple clustering in which each view comprises a co-clustering structure, and each cluster block fits a univariate distribution. Though our method assumes a somewhat complicated cluster structure (multiple views of co-clustering structures), it effectively detects multiple cluster solutions by clustering relevant features within a view, based on their distributional patterns. In contrast with our multiple co-clustering method, the restricted multiple clustering method is simple and straightforward for implementation. However, from a factor-analytical perspective, fitting a single distribution to all features in a view implies the dimensional reduction of that view by a single factor, which may be too restrictive for high-dimensional data. On the other hand, our method relaxes this constraint, allowing possible factors to be inferred in a data-driven approach. Practically, if there is prior knowledge that each view consists of a single factor, then we may use the restricted multiple clustering method. Otherwise, it is preferable to use our multiple co-clustering method, as demonstrated in both synthetic and real data applications above.

In comparisons with COALA and decorrelated K-means methods, our multiple co-clustering method outperforms other state-of-the-art methods using facial image and cardiac arrhythmia data. Beyond its better performance in sample clustering, our multiple co-clustering method has several advantages over other methods. It can infer the number of views/clusters. It is applicable to datasets comprising different types of features, and it can identify relevant features. Furthermore, our method is computationally efficient. The reason for this efficiency is the fitting of a univariate distribution to each cluster block. It is notable that despite using only a univariate distribution, our

multiple co-clustering method can flexibly fit a dataset by adapting the number of views/clusters by means of a Dirichlet process.

Finally, it is worth noting that the multiple co-clustering method is not only useful to recover multiple cluster structures of data, but also a single-cluster structure. In the case of single clustering, our method works by selecting relevant features for possible sample clustering. This may be the main reason that our method performs better with the cardiac arrhythmia data than COALA and decorrelated K means, which use the data without feature selection.

6 Acknowledgement

This study is the result of "Integrated Research on Neuropsychiatric Disorders" carried out under the Strategic Research Program for Brain Science by the Ministry of Education, Culture, Sports, Science and Technology of Japan. We would like to thank Dr. Steve D. Aird at Okinawa Institute of Science and Technology Graduate University for his proof-reading of this article.

Appendices

A Observation models for Gaussian, Poisson and multinomial

This is a supplementary material for Section 2.4 (Observation models) in the main manuscript, providing priors, the expectation of log-likelihood and the updating equations.

A.1 Gaussian distribution

We denote univariate Gaussian density function as $\text{Gauss}(\cdot|\mu, \sigma^2)$ where μ and σ^2 are mean and variances. We assume conjugate priors for μ and σ^2 in each cluster block:

$$\begin{aligned} s_{v,g,k} &\sim \text{Ga}(\cdot|\gamma_0/2, \gamma_0\sigma_0^2/2) \\ \mu_{v,g,k} &\sim \text{Gauss}(\cdot|\mu_0, (\lambda_0 s_{v,g,k})^{-1}), \end{aligned}$$

where $\text{Ga}(\cdot|a, b)$ denotes Gamma distribution with shape and rate parameters (a, b). In the present paper, we set $\sigma_0^2 = 10^4$, $\gamma_0 = 1$, and $\lambda_0 = 10^{-4}$ so that the prior distributions are nearly non-informative. It can be shown that the variational approximation for the posterior $q_{\boldsymbol{\theta}^{(m)}}(\boldsymbol{\theta}^{(m)})$ is given by

$$\begin{aligned} \prod_{v=1}^{V}\prod_{g=1}^{G}\prod_{k=1}^{K} &\text{Gauss}(\mu_{v,g,k}|\mu_{0,v,g,k}, (\lambda_{0,v,g,k}s_{0,v,g,k})^{-1}) \\ &\times \text{Ga}(s_{0,v,g,k}|\gamma_{0,v,g,k}/2, \gamma_{0,v,g,k}\sigma_{0,v,g,k}^2/2), \end{aligned}$$

where the hyperparameters are updated by

$$\begin{aligned}
\lambda_{0,v,g,k} &= \lambda_0 + \sum_{j=1}^{d^{(m)}} \sum_{i=1}^{n} \tau_{j,v,g}^{(m)} \eta_{i,v,k} \\
\mu_{0,v,g,k} &= \frac{1}{\lambda_{0,v,g,k}} \Big\{ \lambda_0 \mu_0 + \sum_{j=1}^{d^{(m)}} \sum_{i=1}^{n} \tau_{j,v,g}^{(m)} \eta_{i,v,k} X_{i,j}^{(m)} \Big\} \\
\gamma_{0,v,g,k} &= \gamma_0 + \sum_{j=1}^{d^{(m)}} \sum_{i=1}^{n} \tau_{j,v,g}^{(m)} \eta_{i,v,k} \\
\sigma_{0,v,g,k}^2 &= \frac{1}{\gamma_{0,v,g,k}} \Big\{ \gamma_0 \sigma_0^2 + \lambda_0 \mu_0^2 \\
&+ \sum_{j=1}^{d^{(m)}} \sum_{i=1}^{n} \tau_{j,v,g}^{(m)} \eta_{i,v,k} (X_{i,j}^{(m)})^2 - \lambda_{0,v,g,k} \mu_{0,v,g,k}^2 \Big\}.
\end{aligned}$$

Finally, the expectation of the conditional log-likelihood $\mathbb{E}_{q(\boldsymbol{\theta})} \Big[\log p(X_{i,j}^{(m)} | \boldsymbol{\theta}_{v,g,k}^{(m)}) \Big]$ is given by

$$\begin{aligned}
-\frac{1}{2} \Big\{ & \frac{(X_{i,j}^{(m)} - \mu_{0,v,g,k})^2}{\sigma_{0,v,g,k}^2} + \frac{1}{\lambda_{0,v,g,k}} + \log \sigma_{0,v,g,k}^2 \\
& + \log(\gamma_{0,v,g,k}/2) - \psi(\gamma_{0,v,g,k}/2) + \log(2\pi) \Big\}.
\end{aligned}$$

A.2 Poisson distribution

We denote Poisson distribution as $\text{Poisson}(\cdot|\lambda)$ where λ is a rate parameter. The conjugate prior for λ is given by

$$\lambda_{v,g,k} \sim \text{Ga}(\cdot|\alpha_0, \beta_0),$$

where we set α_0 and β_0 to one. It can be shown that the variational approximation is given by

$$q_{\boldsymbol{\theta}^{(m)}}(\boldsymbol{\theta}^{(m)}) = \prod_{v=1}^{V} \prod_{g=1}^{G} \prod_{k=1}^{K} \text{Ga}(\lambda_{v,g,k} | \alpha_{0,v,g,k}, \beta_{0,v,g,k}),$$

where the hyperparameters are updated by

$$\begin{aligned}
\alpha_{0,v,g,k} &= \alpha_0 + \sum_{j=1}^{d^{(m)}} \sum_{i=1}^{n} \tau_{j,v,g}^{(m)} \eta_{i,v,k} X_{i,j}^{(m)} \\
\beta_{0,v,g,k} &= \beta_0 + \sum_{j=1}^{d^{(m)}} \sum_{i=1}^{n} \tau_{j,v,g}^{(m)} \eta_{i,v,k}.
\end{aligned}$$

The expectation of the conditional log-likelihood becomes

$$X_{i,j}^{(m)}\{\psi(\alpha_{0,v,g,k}) - \psi(\beta_{0,v,g,k})\} - \frac{\alpha_{0,v,g,k}}{\beta_{0,v,g,k}} - \sum_{t=1}^{X_{i,j}^{(m)}} \log t.$$

A.3 Categorical/multinomial distribution

For a categorical feature x ($x \in \{c_1, \ldots, c_H\}$), we denote categorical distribution as $\text{Cat}(\cdot|\boldsymbol{p})$ where H is the number of categories, and $\boldsymbol{p} = (p_1, \ldots, p_H)$ are probabilities for each category with $\sum_{h=1}^{H} p_h = 1$. We assume the conjugate prior for $(p_1, \ldots, p_H)$,

$$(p_1, \ldots, p_H) \quad \sim \quad \text{Dirichlet}(\cdot|\boldsymbol{\rho}_0),$$

where $\text{Dirichlet}(\cdot|\boldsymbol{\rho}_0)$ denotes a Dirichlet distribution with prior sample size $\boldsymbol{\rho}_0$. We set $\boldsymbol{\rho}_0$ to $(1, \ldots, 1)$. It can be shown that

$$q_{\boldsymbol{\theta}^{(m)}}(\boldsymbol{\theta}^{(m)}) = \prod_{v=1}^{V}\prod_{g=1}^{G}\prod_{k=1}^{K} \text{Dirichlet}(\boldsymbol{p}_{v,g,k}|\boldsymbol{\rho}_{0,v,g,k}),$$

where the hyperparameters are updated by

$$\rho_{0,v,g,k,h} = \rho_{0,h} + \sum_{j=1}^{d^{(m)}}\sum_{i=1}^{n} \tau_{j,v,g}^{(m)}\eta_{i,v,k}\mathbb{I}(X_{i,j}^{(m)} = c_h),$$

where $\rho_{0,v,g,k,h}$ denotes the hth element of $\boldsymbol{\rho}_{0,v,g,k}$. The expectation of the log-likelihood is then given by

$$\sum_{h=1}^{H} \mathbb{I}(X_{i,j}^{(m)} = c_h)\{\psi(\rho_{0,h,v,g,k}) - \psi(\sum_{h'=1}^{H} \rho_{0,h',v,g,k})\}.$$

Since the categorical distribution differs depending on the number of categories H, we need to define different types of categorical distribution. Alternatively, for the purpose of simplicity, we can set H to the maximum number of categories for different categorical features, and fit a single family of categorical distribution to all these features.

More generally, in the case of multinomial distribution, the update equation and the expectation of the log-likelihood becomes

$$\rho_{0,v,g,k,h} = \rho_{0,h} + \sum_{j=1}^{d^{(m)}}\sum_{i=1}^{n} \tau^{(m)}\eta_{i,v,k}n_{i,j,h}$$

$$\sum_{h=1}^{H} n_{i,j,h}\{\psi(\rho_{0,h,v,g,k}) - \psi(\sum_{h'=1}^{H} \rho_{0,h',v,g,k})\}$$

$$+ \log \binom{\sum_{h=1}^{H} n_{i,j,h}}{n_{i,j,1}, \ldots, n_{i,j,H}},$$

where $n_{i,j,h}$ is the number of category c_h in $X_{i,j}^{(m)}$; the last term is the logarithm of multinomial coefficients.

B List of features for clinical data

Table 6: *List of features for clinical data*

Features
Numerical features
BAS (Behavioral Activation Scale),
BDNF (Quantity of brain-derived neurotrophic factor in blood),
BDI2 (Beck Depression Inventory),
BIS (Behavioral Inhibition Scale),
CATS (Child Abuse and Trauma Scale),
Cortisol (Quantity of cortisol in blood),
GAF (Global Assessment of Functioning),
PHQ9 (Patient Health Questionnaire),
HRSD17 (Hamilton Rating Scale for Depression),
JART (Adult reading test),
LES (Life Experiences Survey),
PANASP (Positive Affect Schedule),
PANASN (Negative Affect Schedule),
SHAPS (Snaith-Hamilton Pleasure Scale),
STAI (State-Trait Anxiety Inventory),
N, E, O, A, C
(Five factors in revised NEO Personality Inventory)
Categorical features
Sex
miniA-P (Mini-International Neuropsychiatric Interview),
A-P corresponds to the following psychiatric symptoms:
Major depressive disorder (A),
Dysthymia (B), Suicide risk (C),
Mania (D), Panic disorder (E), Agoraphobia (F),
Social phobia (G),
Obsessive compulsive disorder (H),
PTSD (I), Alcohol dependence and abuse (J),
Drug dependence and abuse (K),
Psychotic disorder (L), Anorexia (M), Bulimia (N),
Generalized anxiety disorder (O),
Antisocial personality disorder (P),
SNPs 1-8: Single Nucleotide Polymorphisms that are located in the following genome sites, respectively. (in parenthesis are the relevant gene functions)
rs1187323 (NTRK2), rs34118353 (5HT1a receptor),
rs3756318 (NTRK2), rs3813929 (5HT2c receptor),
rs45554739 (NTRK2), rs56384968 (SLC6A4),
rs6265 (BDNF), rs6294 (5HT1a receptor)

References

[1] Eric Bae and James Bailey. Coala: A novel approach for the extraction of an alternate clustering of high quality and high dissimilarity. In *Data Mining, 2006. ICDM'06. Sixth International Conference on*, pages 53–62. IEEE, 2006.

[2] David M Blei, Michael I Jordan, et al. Variational inference for dirichlet process mixtures. *Bayesian analysis*, 1(1):121–143, 2006.

[3] Rich Caruana, M Elhaway, Nam Nguyen, and Casey Smith. Meta clustering. In *Data Mining, 2006. ICDM'06. Sixth International Conference on*, pages 107–118. IEEE, 2006.

[4] Xuan Hong Dang and James Bailey. Generation of alternative clusterings using the cami approach. In *SDM*, volume 10, pages 118–129. SIAM, 2010.

[5] Michael D Escobar and Mike West. Bayesian density estimation and inference using mixtures. *Journal of the american statistical association*, 90(430):577–588, 1995.

[6] Thomas S Ferguson. A bayesian analysis of some nonparametric problems. *The annals of statistics*, pages 209–230, 1973.

[7] Andrew Gelman, John B Carlin, Hal S Stern, and Donald B Rubin. *Bayesian data analysis*, volume 2. Taylor & Francis, 2014.

[8] Yue Guan, Jennifer G Dy, Donglin Niu, and Zoubin Ghahramani. Variational inference for nonparametric multiple clustering. In *MultiClust Workshop, KDD-2010*, 2010.

[9] H Güvenir, Burak Acar, Gülsen Demiröz, et al. A supervised machine learning algorithm for arrhythmia analysis. In *Computers in Cardiology 1997*, pages 433–436. IEEE, 1997.

[10] Lawrence Hubert and Phipps Arabie. Comparing partitions. *Journal of classification*, 2(1):193–218, 1985.

[11] Prateek Jain, Raghu Meka, and Inderjit S Dhillon. Simultaneous unsupervised learning of disparate clusterings. *Statistical Analysis and Data Mining: The ASA Data Science Journal*, 1(3):195–210, 2008.

[12] Valdemar Jensen. Sur les fonctions convexes et les inégalités entre les valeurs moyennes. *Acta Mathematica*, 30(1):175–193, 1906.

[13] Emmanuel Muller, Stephan Gunnemann, Ines Farber, and Thomas Seidl. Discovering multiple clustering solutions: Grouping objects in different views of the data. In *Data Engineering (ICDE), 2012 IEEE 28th International Conference*, pages 1207–1210. IEEE, 2012.

[14] Kevin Murphy. *Machine Learning: A Probabilistic Perspective*. Cambridge, Massachusetts: MIT Press, 2012.

[15] Donglin Niu, Jennifer G Dy, and Zoubin Ghahramani. A nonparametric bayesian model for multiple clustering with overlapping feature views. In *International Conference on Artificial Intelligence and Statistics*, pages 814–822, 2012.

[16] Donglin Niu, Jennifer G Dy, and Michael I Jordan. Multiple non-redundant spectral clustering views. In *Proceedings of the 27th international conference on machine learning (ICML-10)*, pages 831–838, 2010.

[17] Hanhuai Shan and Arindam Banerjee. Bayesian co-clustering. In *Data Mining, 2008. ICDM'08. Eighth IEEE International Conference*, pages 530–539. IEEE, 2008.